
On Sampling and Greedy MAP Inference of Constrained Determinantal Point Processes

Tarun Kathuria, Amit Deshpande

Microsoft Research, India

t-takat@microsoft.com, amitdesh@microsoft.com

Abstract

Subset selection problems ask for a small, diverse yet representative subset of the given data. When pairwise similarities are captured by a kernel, the determinants of submatrices provide a measure of diversity or independence of items within a subset. Matroid theory gives another notion of independence, thus giving rise to optimization and sampling questions about Determinantal Point Processes (DPPs) under matroid constraints. Partition constraints, as a special case, arise naturally when incorporating additional labeling or clustering information, besides the kernel, in DPPs. Finding the maximum determinant submatrix under matroid constraints on its row/column indices has been previously studied. However, the corresponding question of sampling from DPPs under matroid constraints has been unresolved, beyond the simple cardinality constrained k -DPPs.

We give the first polynomial time algorithm to sample exactly from DPPs under partition constraints, for any constant number of partitions. We complement this by a complexity theoretic barrier that rules out such a result under general matroid constraints. Our experiments indicate that partition-constrained DPPs offer more flexibility and more diversity than k -DPPs and their naive extensions, while being reasonably efficient in running time.

We also show that a simple greedy initialization followed by local search gives improved approximation guarantees for the problem of MAP inference from k -DPPs on *well-conditioned* kernels. Our experiments show that this improvement is significant for larger values of k , supporting our theoretical result.

1 Introduction

Selecting a small, diverse yet representative subset of the given data is an important problem underlying feature/exemplar selection in machine learning, sensor placements [22], row/column subset selection in linear algebra [14], and coresets in computational geometry [13]. Subset selection from a given *ground set* \mathcal{Y} of m items is often formulated as an optimization or a sampling problem over all the 2^m possible subsets, with constraints on the size and the diversity of the desired subset. If an m -by- m positive semidefinite kernel matrix K captures all pairwise similarities, then matrix determinants and log-determinants help formulate diversity in the form of volume [14, 15, 19], entropy [22], and repulsion of Fermions [32]. Geometrically, this corresponds to picking a subset of vectors that are long and as linearly independent as possible. Matroid theory provides another compelling notion of independence in subsets that underlies spanning trees, matchings etc. Thus, a large class of subset selection problems can be formulated as either optimization or sampling problems, stated using determinants or log-determinants, and subject to various matroid constraints.

Motivated by the above, we model the diversity of a subset $S \subseteq [m]$ by the determinant of its corresponding $|S|$ -by- $|S|$ submatrix $K_{S,S}$ in a given m -by- m positive semidefinite kernel K , and summarize the previous work on optimization and sampling questions to lay the premise.

The optimization problem for subset selection corresponds to maximizing $\det(K_{S,S})$, or equivalently $\log \det(K_{S,S})$, subject to additional cardinality or matroid constraints. This is NP-hard even in the absence of constraints [10]. However, the submodularity of $f(S) = \log \det(K_{S,S})$ gives a polynomial time $1/4$ -approximation for maximizing $\log \det(K_{S,S})$ [18]. When $\log \det(K_{S,S})$ is monotone (e.g., when the smallest eigenvalue $\lambda_{\min}(K) \geq 1$), a greedy algorithm gives $(1 - 1/e)$ -approximation, even under cardinality or matroid constraints [29]. However, multiplicative approximations for maximizing $\log \det(K_{S,S})$ do not imply similar results for $\det(K_{S,S})$. The best known polynomial time algorithm for maximizing $\det(K_{S,S})$ over subsets of size k gives an approximation guarantee of $1/e^{k+o(k)}$ [30], and it is NP-hard to do better than $1/c^k$, for some $c > 1$ [33]. Recently, the same approximation guarantee was achieved for maximizing $\det(K_{S,S})$ under partition constraints on S , which is a special case of matroid constraints, via a *geometric concave program* [31].

On the other hand, the sampling problem for subset selection corresponds to sampling a subset $S \subseteq [m]$ with probability proportional to $\det(K_{S,S})$. These distributions are known as Fermion Point Processes or Determinantal Point Processes (DPPs), with remarkable applications in probability, statistical physics, and random matrix theory [6, 7, 28, 20]. Sampling exactly from a DPP and its cardinality-constrained variant k -DPP can both be done in polynomial time [14, 20]. This has found applications in document summarization [26], object retrieval [2], sensor placement [22]. We point the reader to [23, 24] for efficient algorithmic sampling, inference, and computation of the partition functions, marginals, conditional probabilities etc. of DPPs and k -DPPs, and their numerous applications in machine learning.

Beyond the simple cardinality-constrained k -DPPs, previous work does not provide efficient algorithms for exact sampling from a DPP under general matroid constraints. The difficulty arises from the computation of its partition function, namely, the sum of $\det(K_{S,S})$ over subsets S that satisfy the given matroid constraints. For general matroid constraints, this sum does not have a nice, closed form expression (unlike in the case of DPPs and k -DPPs).

The main purpose of our paper is to study sampling from DPPs under general matroid constraints. We give the first polynomial time algorithm for exact sampling from DPPs under partition constraints, for any constant number of partitions. We complement this by a complexity theoretic barrier that rules out such a result under general matroid constraints. To be precise, DPPs under transversal matroid constraints can simulate random perfect matchings in a bipartite graph, whose partition function is $\#P$ -hard (refer to Appendix 1 for details).

Partition constraints arise naturally when we have some inherent labeling or clustering of the items that is not captured by the DPP kernel. In such cases, DPPs may not capture the diversity correctly, leading to over/under-representation as explained below.

1. Google image search for the query “jaguar” fetches images of the animal, the cars from the company “Jaguar”, the logo of the car company etc. While the search engine recognizes these categories, the top results are dominated by the animal and the cars, causing the company logo to be under-represented. k -DPPs based only on image kernels are unlikely to fix this problem as the stylistic features of the logo are similar to that of the animal.
2. Consider a facial image database of people with different expressions and lighting conditions. Kernels based on SIFT and dot-products do not capture lighting conditions well but are still desirable due to their computational efficiency. Now if we want a subset of facial images where the people as well as the lighting conditions therein are distinct, or if want certain lighting conditions more prominently in our subset, then the kernel alone may not be able to capture such constraints.

In these two examples, partition constraints arise naturally from the image categories and the lighting conditions, respectively. They could also arise via clustering based on features not captured by the kernel. Thus, when the ground set has a disjoint partition $\mathcal{V} = \mathcal{P}_1 \uplus \mathcal{P}_2 \uplus \dots \uplus \mathcal{P}_p$ given apriori, and we want to pick a subset S with k_1 items from \mathcal{P}_1 , k_2 items from \mathcal{P}_2 , and so on, we define a Partition-DPP as the corresponding conditioning of DPP under these constraints on S .

It is imperative to contrast Partition-DPP with two natural extensions of k -DPP, namely, (a) sampling $k = k_1 + \dots + k_p$ images from the entire set by k -DPP and (b) sampling from p independent k_i -DPPs each on part \mathcal{P}_i , respectively. k -DPP suffers from over/under-representation and is not guaranteed to pick k_i items from \mathcal{P}_i , due to interference between inter-partition and intra-partition diversity.

In the example of different people under different lighting conditions, if our partitions are given by the lighting conditions and we sample from independent k_i -DPPs on each part, we may end up picking the same person from different lighting conditions in our subset. It is also easy to construct counterexamples to rule out k -DPP followed by rejection sampling based on the partition constraints.

Our algorithm for Partition-DPPs is based on a multivariate generalization of the characteristic polynomial of a matrix that may be of independent interest. Since the effectiveness of DPPs in modeling diversity has been demonstrated in [23, 24], our experiments focus on comparing our algorithm against the two natural extensions of k -DPP stated above, and highlight the over/under-representation, or equivalently, the violation of partition constraints.

In addition, we quantitatively improve the MAP inference approximation guarantee for k -DPPs, when the DPP kernel is *well-conditioned*, and show that a simple greedy algorithm followed by local search provides almost as good an approximation guarantee for maximizing $\det(K_{S,S})$ over k -sized subsets as the expensive convex program in [30]. Our experiments show that this improvement is significant for larger values of k , supporting our theoretical result.

2 Setup and basic definitions

Let $[m]$ denote the set $\{1, 2, \dots, m\}$ used to index the m items in our ground set \mathcal{Y} . We identify \mathcal{Y} with $[m]$ and may use them interchangeably. For any positive semidefinite matrix $K \in \mathbb{R}^{m \times m}$ and a subset $S \subseteq [m]$, we use $K_{S,S}$ to denote the corresponding $|S|$ -by- $|S|$ submatrix of K whose both the row and the column indices are in S . For a rectangular matrix $A \in \mathbb{R}^{m \times n}$, we denote its rows by $a_1, a_2, \dots, a_m \in \mathbb{R}^n$, and we use A_S to denote the $|S|$ -by- n row-submatrix formed by row indices in S . Let $\text{span}(S)$ be the linear subspace spanned by the rows of A_S , and $\pi_S(A)$ be the orthogonal projection onto $\text{span}(S)$.

For $K = AA^T$, we have $K_{S,S} = A_S A_S^T$. k -DPP defines a distribution over subsets $S \subseteq [m]$ of size k , where the probability of picking S is proportional to $\det(K_{S,S}) = \det(A_S A_S^T)$, which is also proportional to the squared volume of the parallelepiped formed by the rows of A_S .

Definition 2.1. (k -DPP) Given a positive semidefinite matrix $K \in \mathbb{R}^{m \times m}$ as DPP kernel, the cardinality-constrained DPP k -DPP is defined by the following distribution on subsets $B \subseteq [m]$ of size k .

$$\Pr(B) = \frac{\det(K_{B,B})}{\sum_{S: |S|=k} \det(K_{S,S})} = \frac{\det(A_B A_B^T)}{\sum_{S: |S|=k} \det(A_S A_S^T)}.$$

The following linear algebraic identity about the coefficients of the characteristic polynomial is at the heart of the polynomial time sampling algorithm for k -DPPs in [14, 23] (see Algorithm 1).

Theorem 2.2. (Proposition 3.2 in [15]) For any $A \in \mathbb{R}^{m \times n}$, let the characteristic polynomial of $AA^T \in \mathbb{R}^{m \times m}$ be $\det(AA^T - xI) = x^m + c_{m-1}x^{m-1} + \dots + c_0$. Then

$$\sum_{S: |S|=k} \det(A_S A_S^T) = |c_{m-k}|, \quad \text{for } 1 \leq k \leq m.$$

We refer the reader to [15, 3] for the proof of the above theorem.

Given a partition of the ground set into disjoint parts, a partition constraint on S says that it must have k_1 items from the first part, k_2 from the second part, and so on.

Definition 2.3. (Partition Constraint) Given a disjoint partition \mathcal{P} as $\mathcal{Y} = [m] = \mathcal{P}_1 \uplus \mathcal{P}_2 \uplus \dots \uplus \mathcal{P}_p$ into p parts of sizes m_1, m_2, \dots, m_p , respectively, a partition constraint on a subset $S \subseteq [m]$ is defined by a p -tuple (k_1, k_2, \dots, k_p) , and S is said to satisfy this partition constraint if S has k_i elements from part \mathcal{P}_i , respectively, i.e., $|S \cap \mathcal{P}_i| = k_i$, for $1 \leq i \leq p$.

Now we define a multivariate characteristic polynomial of AA^T with respect to the partition $\mathcal{Y} = [m] = \mathcal{P}_1 \uplus \mathcal{P}_2 \uplus \dots \uplus \mathcal{P}_p$, by splitting xI in $\det(AA^T - xI)$ into smaller identity matrices with different variables, one for each partition.

Definition 2.4. (Multivariate Characteristic Polynomial) For any $A \in \mathbb{R}^{m \times n}$ and a partition \mathcal{P} as $\mathcal{Y} = [m] = \mathcal{P}_1 \uplus \mathcal{P}_2 \uplus \dots \uplus \mathcal{P}_p$, we define the corresponding multivariate characteristic polynomial

Algorithm 1 Sampling from k -DPP (Algorithm 1 in [14])

Input: Matrix $A \in \mathbb{R}^{m \times n}$ and $1 \leq k \leq \text{rank}(A)$
Output: Subset S of k rows of A picked with probability proportional to $\det(A_S A_S^T)$
Initialize $S \leftarrow \emptyset$ and $B \leftarrow A$.
for $t = 1$ **to** k **do**
 for $i = 1$ **to** m **do**
 Compute $p_i = \|b_i\|^2 |c_{m-k+t}(C_i C_i^T)|$,
 where C_i is the matrix obtained by projecting each row of B orthogonal to b_i .
 end for
 Pick i with probability proportional to p_i .
 $S \leftarrow S \cup \{i\}$ and $B \leftarrow C_i$
end for
return S

of AA^T as

$$\det(AA^T - x_1 I_1 - \dots - x_p I_p) = \sum_{i_p=0}^{m_p} \dots \sum_{i_1=0}^{m_1} c_{i_1, \dots, i_p} x_1^{i_1} \dots x_p^{i_p}, \text{ where } (I_t)_{i,i} = \begin{cases} 1 & \text{if } i \in \mathcal{P}_t \\ 0 & \text{otherwise} \end{cases}.$$

3 Main Results

In this section, we first present our sampling algorithm for Partition-DPPs that uses the multivariate characteristic polynomial to generalize k -DPP sampling. In the later subsection, we present our results about quantitative improvements for MAP inference in k -DPPs, for *well-conditioned* DPP kernels, using greedy algorithm followed by local search.

3.1 Partition-DPPs

As argued earlier, k -DPPs and independent k_i -DPPs for each partition cannot be used for exact sampling from Partition-DPPs. It is also easy to see that k -DPPs followed by rejection sampling based on the partition constraints cannot give an efficient algorithm either. In [21] (Section 3.6), the authors consider Submodular Point Processes and Log-submodular Point Processes (of which DPPs are a special case) under partition constraints of the type $|S \cap \mathcal{P}_i| \leq k_i$, for $1 \leq i \leq p$. They argue that the partition function Z , that is, the sum of $\det(K_{S,S})$ over subsets S satisfying such partition constraints, splits as $Z = \sum_{l=1}^p \sum_{k=1}^{k_l} Z_{\mathcal{P}_l}^k$, where each $Z_{\mathcal{P}_l}^k$ is the partition function of a k -DPP defined only on the elements of \mathcal{P}_l . This partition function Z actually corresponds a mixture of cardinality-constrained DPPs on each part, and does not apply to Partition-DPP as in our definition [5]. We explain this with an example in the appendix.

Now we show how the coefficients of the multivariate characteristic polynomial are related to the partition functions for Partition-DPPs, that is, the sum of $\det(K_{S,S})$ over all subsets S that satisfy a given partition constraint.

Lemma 3.1. *Consider the multivariate characteristic polynomial of $AA^T \in \mathbb{R}^{m \times m}$ with respect to partition \mathcal{P} given by $\mathcal{Y} = \mathcal{P}_1 \uplus \mathcal{P}_2 \uplus \dots \uplus \mathcal{P}_p$, as in Definition 2.4. Then*

$$\sum_{S : |S \cap \mathcal{P}_i| = k_i, \text{ for } 1 \leq i \leq p} \det(A_S A_S^T) = |c_{m_1-k_1, \dots, m_p-k_p}|, \quad \text{for any } 0 \leq k_i \leq m_i.$$

Proof. (Sketch.) We first expand $\det(AA^T - x_1 I_1 - \dots - x_p I_p)$ using the Leibniz formula for determinant [34], and then collect the terms for each monomial together. Again by Leibniz formula, the collected terms sum up $\det(A_S A_S^T)$ over certain subsets S , which happen to be the ones that satisfy the corresponding partition constraints. \square

Algorithm 2 Sampling from Partition-DPPs

Input: Matrix $A \in \mathbb{R}^{m \times n}$ and $1 \leq k \leq \text{rank}(A)$
Output: Subset S of k rows of A picked with probability proportional to $\det(A_S A_S^T)$
Initialize $S \leftarrow \emptyset$ and $B \leftarrow A$.
for $t = 0$ **to** $k - 1$ **do**
 for $i = 1$ **to** m **do**
 Compute $p_i = \|b_i\|^2 |c'(C_i C_i^T)|$,
 where C_i is a matrix obtained by projecting each row of B orthogonal to b_i
 and $c'(C_i C_i^T)$ is defined as per Theorem 3.2.
 end for
 Pick i with probability proportional to p_i .
 $S \leftarrow S \cup \{i\}$ and $B \leftarrow C_i$
end for
return S

As in the case of k -DPPs or volume sampling [14], it helps to think of Partition-DPP as a distribution over ordered tuples rather than sets. Specifically, $\Pr(X_1 = i_1, \dots, X_k = i_k)$ is equal to

$$\begin{cases} \frac{\det(A_{\{i_1, \dots, i_k\}} A_{\{i_1, \dots, i_k\}}^T)}{k! \sum_{S: |S \cap \mathcal{P}_i| = k_i, \text{ for } 1 \leq i \leq p} \det(A_S A_S^T)}, & \text{for distinct } i_1, \dots, i_k \text{ that} \\ & \text{satisfy partition constraints } \mathcal{P} \\ 0, & \text{otherwise} \end{cases}$$

Then the marginal probabilities $\Pr(X_{t+1} = i \mid X_1 = i_1, \dots, X_t = i_t)$ have the following interpretation in terms of the co-efficients of certain derived multivariate characteristic polynomials. This generalizes a corresponding theorem for k -DPP in [14] and gives the technical core of our algorithm.

Theorem 3.2. *Let $(i_1, \dots, i_t) \in [m]^t$ be such that $\Pr(X_1 = i_1, \dots, X_t = i_t) > 0$, in the above Partition-DPP extended to k -tuples. Let $S = \{i_1, \dots, i_t\}$ and $|S \cap \mathcal{P}_l| = t_l$, for all $l \in [p]$, with $t = \sum_{l=1}^p t_l$. For any $i \in \mathcal{P}_j$, let $B = A - \pi_S(A)$ and $C_i = B - \pi_{\{i\}}(B) = A - \pi_{S \cup \{i\}}(A)$. Then*

$$\Pr(X_{t+1} = i \mid X_1 = i_1, \dots, X_t = i_t) = \frac{\|b_i\|^2 |c'(C_i C_i^T)|}{(k - t) |c''(B B^T)|}, \quad \text{where}$$

$$c'(C_i C_i^T) = c_{m_1 - k_1 + t_1, \dots, m_j - k_j + t_j + 1, \dots, m_p - k_p + t_p}(C_i C_i^T)$$

$$c''(B B^T) = c_{m_1 - k_1 + t_1, \dots, m_j - k_j + t_j, \dots, m_p - k_p + t_p}(B B^T).$$

Proof. (Sketch.) Being conditional probability, $\Pr(X_{t+1} = i \mid X_1 = i_1, \dots, X_t = i_t)$ can be written as $\Pr(X_1 = i_1, \dots, X_t = i_t, X_{t+1} = i) / \Pr(X_1 = i_1, \dots, X_t = i_t)$, where the numerator and the denominator sum up probabilities over the indices in (i_1, \dots, i_k) not fixed by the marginal probability expression. Thus, the numerator sums over subsets of the remaining indices that satisfy partition constraints of the form, $k_1 - t_1, \dots, k_j - (t_j + 1), \dots, k_p - t_p$. The denominator is a similar sum over subsets with partition constraints of the form $k_1 - t_1, \dots, k_j - t_j, \dots, k_p - t_p$. Some further manipulations of the subdeterminants in the sums and Lemma 3.1 complete the proof. \square

With this theorem in hand, Algorithm 2 follows naturally. By computing the co-efficients of multivariate characteristic polynomials as in Definition 2.4, we can do exact sampling from DPPs under partition constraints. Algorithm 2 runs in time $\mathcal{O}(m^{p+4}nk/p^p)$, where $k = \sum_{l=1}^p k_l$. One way to compute the coefficients of a univariate characteristic polynomial is by evaluating the polynomial $\det(A - xI)$ at $m + 1$ distinct points x , which gives a system of linear equations in its coefficients that can be solved by inverting a Vandermonde matrix. This generalizes to the multivariate case, albeit inefficiently, because it can have $\prod_{l=1}^p (m_l + 1) \leq (1 + m/p)^p$ coefficients (by AM-GM inequality). Also this interpolation approach computes all the coefficients, when the algorithm requires only one of them to proceed. This causes the exponential in p blowup in our running time at present.

3.2 MAP Inference for k -DPPs

The maximum a posteriori (MAP) inference problem for k -DPPs corresponds to finding a subset S of size k that maximizes $\det(A_S A_S^T)$. This problem is NP-hard to approximate within a multiplicative factor of $2^{\mathcal{O}(k)}$ [12], a simple greedy algorithm gives a $2^{\mathcal{O}(k \log k)}$ -approximation, and the only improvement over the greedy that achieves a $2^{\mathcal{O}(k)}$ -approximation is via a convex program that is practically prohibitive [30]. Thus, a natural question is whether a greedy initialization can be further improved by some local search heuristic. In this subsection, we give a simple $2^{k \log \min\{k, \kappa(A)\}}$ -approximation algorithm this way. Given an $\epsilon > 0$, start with the greedy solution S , and pick $i \in S$ and $j \notin S$ such swapping i and j gives the maximum improvement over the current solution S . We do this swap only if the improvement is at least $(1 + \epsilon/k)$ in multiplicative factor. Our experiments show that this gives significantly better results as k grows larger.

Theorem 3.3. *A greedy initialization followed by $\mathcal{O}(k^2 \log k / \epsilon)$ local search steps as above gives a $2^{k \log \min\{k, \kappa(A)\}}$ -approximation to $\max_{|S|=k} \det(A_S A_S^T)$, where $\kappa(A)$ is*

$$\kappa(A) = \frac{\lambda_1(AA^T)}{\frac{1}{n-k+1} \sum_{t \geq k} \lambda_t(AA^T)}.$$

Note that $\kappa(A)$ is at most the condition number of AA^T but in practice could be much better if the ill-conditioning is only due to a small number of bottom eigenvalues. This algorithm takes time linear in the number of non-zero entries of A .

Proof. See the Appendix. In short, the proof relates $\det(A_S A_S^T)$ to the rank- k approximation error for AA^T , which is why we get $\kappa(A)$ defined as above. \square

4 Experiments

The effectiveness of DPPs in modeling diversity has been demonstrated in [23, 24]. Our experiments focus on comparing Partition-DPP algorithm against the two natural extensions of k -DPP to highlight their over/under-representation and other issues in a way that makes them evident to the reader.

4.1 Limitations of k -DPPs and independent k_i -DPPs

We study the performance of Partition-DPPs on two real-world image search tasks as follows.

1. Our first task compares k -DPPs, where $k = \sum_{i=1}^p k_i$, against DPPs that take partition into account, namely, the independent k_i -DPPs on each partition and the Partition-DPP. We compare the skew in the number of items per part returned by k -DPP to measure its over/under-representation. Note that both the independent k_i -DPPs and Partition-DPPs pick exactly k_i items from part i .
2. Our second task highlights a disadvantage of independent k_i -DPPs on multi-label images. Suppose each label forms a partition. To ensure disjoint partitions, an image is replicated for its each label and each copy is treated as a separate item. Now independent k_i -DPPs may pick the same image multiple times, coming from different partitions, whereas Partition-DPP does not. Due to lack of space, details of this experiment are provided in the appendix.

4.1.1 Data and Methods

For the first task, we consider the Yale face database [17, 25] with the cropped images under different lighting conditions. We randomly selected a subset of size 32 out of 38 people under two different poses. For each pose, we selected three lighting conditions - light, medium and dark, which form the partitions. We split the people into four groups of size 8 each. In each group, we define two experiments based on the facial expression. We used the cropped dataset since the images are well-centered, aligned and of the same size. Thus, we create and run 8 different experiments. We computed the Scalable Invariant Feature Transform (SIFT) descriptors [27] and the corresponding similarity scores between images using the `vlfeat` toolbox for the first task.

As mentioned above, we compare the Partition-DPP against the k -DPP and the independent k_i -DPP. For the first task, where we have three partitions, we always selected 6 images. For the second task, we are interested in demonstrating the difference between independent k_i -DPPs and Partition-DPPs which stems from the possibility of the former choosing the same element multiple times across different partitions. The 3-tuples considered in this case are $(3, 3, 3)$ and $(2, 3, 2)$.

4.1.2 Results: Yale Face Database Task

This task was performed to illustrate that k -DPPs can end up giving arbitrarily skewed results. Table 1 shows the various elements that were chosen by the 6-DPP ($k = 6$) with the columns separating the elements based on the lighting conditions. People have been labelled from 1 to 8. As can be

| Light | Medium | Dark | Light | Medium | Dark |
|---------|--------|------|-------|--------|------|
| 2,3,5,6 | 7 | 8 | 4,5,8 | 6,7 | 3 |
| 1,3,4 | 6 | 1,5 | 2,5,7 | 1 | 4,6 |
| 3,5 | 2,4,7 | 6 | 1,4,5 | 3,7 | 8 |
| 4,6,8 | 1 | 2,3 | 1,5 | 2,7 | 3,4 |

Table 1: Elements sampled from 6-DPP for the 8 different Yale Task Experiments

seen, the number of elements in a single partition may be quite high/low or be evenly spread out, although it seems that the SIFT kernel was promoting more well-lit images as compared to the other lighting conditions. This inherent bias of the kernel may be undesirable in presence of partition constraints, e.g. when we want to promote medium or dark images more. Independent k_i -DPPs and Partition-DPPs provide the freedom to decide how skewed or uniform the three partitions should be. It should be noted that there were some cases where the same person ended up getting picked due to the SIFT kernel ended up treating the well-lit and dark images of a single person as different. We also show the results obtained by all three methods for the setting of $(k_1, k_2, k_3) = (2, 2, 2)$ in Table 2.



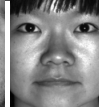




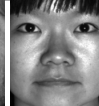

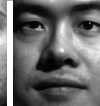



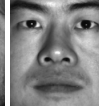




| Method | Light | | | | Medium | | Dark | |
|------------------|---|---|---|---|--|--|---|---|
| 6-DPP |  |  |  |  |  | |  | |
| $(2, 2, 2)$ -DPP | |  |  | |  |  |  |  |
| Partition-DPP | |  |  | |  |  |  |  |


Table 2: Sets obtained by sampling from 6-DPP, $(2,2,2)$ -DPP and the Partition-DPP

4.2 Timing Experiments for Partition-DPPs


While the sampling algorithm for Partition-DPPs is exponential in the number of parts, we present some experiments on synthetic data (since sampling time is independent of kernel and data) with a small number of parts for k -DPPs, k_i -DPPs and Partition-DPPs. The running time is in seconds. In each table, the first row denotes the number of items in each part (m_1, \dots, m_p) and the number of items to be picked from each part (k_1, \dots, k_p) , respectively. The number of parts p will be clear from context. The next three rows show the method names followed by the running time for that method. As seen in the table, our algorithm takes some time when there are roughly 50 elements in total and the number of partitions is 3 and beyond. For $p = 4$, our algorithm takes far too long when m is around 50. However, given that prior work [18, 23] with DPPs typically has m in the range of [50, 200], our algorithm performs reasonably.

| | [15,15], [3,3] | [15,15], [5,5] | [24,24], [5,5] | [24,24], [10,10] |
|---------------|---------------------|---------------------|---------------------|----------------------|
| k -DPP | 0.51 | 0.85 | 2.31 | 5.17 |
| k_i -DPPs | 0.58 | 1.01 | 2.92 | 5.77 |
| Partition-DPP | 3.18 | 7.40 | 79.45 | 133.13 |
| | [10,10,10], [2,2,2] | [10,10,10], [5,5,5] | [16,16,16], [5,5,5] | [6,6,6,6], [1,1,1,1] |
| k -DPP | 0.51 | 0.85 | 2.31 | 5.17 |
| k_i -DPPs | 0.58 | 1.01 | 2.92 | 5.77 |
| Partition-DPP | 3.18 | 7.40 | 79.45 | 133.13 |

Table 3: Results for timing experiments for Partition-DPPs. Time is in seconds.



(a) Synthetic Dataset



(b) Matched Summarization Dataset

Figure 1: Results for MAP Inference for k -DPPs

4.3 MAP Inference for k -DPPs

We compare against the standard Greedy algorithm and the unconstrained Softmax Maximization approach by [18]. For the number of elements k chosen by the Symmetric Greedy algorithm [9] which gives a $1/3$ -approximation ratio for the unconstrained $\log \det$ maximization problem, we compare Greedy and our method as well.

Our first synthetic dataset consists of 50 points in 40 dimensions and each $A_{i,j} \sim 5 * \text{Uniform}(0, 1)$. To recall, $K = AA^T$. In this case, Symmetric greedy picked 39 elements. The results are shown in Figure 1a. Initially, our algorithm does not give any improvement over the Greedy algorithm but as k grows beyond 20, our algorithm starts giving better results. This is in line with our theoretical approximation ratio as well since initially k is smaller than κ but for large k , if the growth of κ is sublinear, our algorithm starts doing better. We refer the reader to the appendix for exact values for this and another synthetically generated dataset.

We also present our results on a real-world dataset used by [18] for matched summarization. We describe our experiment, which is slightly different than their task of selecting document pairs which have additional similarity measures. Statements of eight main contenders in the 2012 US Republican primary debates are extracted and the same preprocessing as [18] is performed. Then, a feature matrix W where W_{qt} is the number of times term t appears in quote q . The kernel considered is $K = WW^T$. We chose a total of 400 quotes (50 from each candidate). Symmetric Greedy picked 312 elements. The results are presented in Figure 1b.

Conclusions and Future Work

We introduce Partition-DPPs for subset selection that respects given partitions or labels that may not be captured by the DPP kernel. Our experiments demonstrate how naive extensions to k -DPPs suffer from over/under-representation or duplication, both of are avoided by Partition-DPPs. We quantitatively improve MAP inference for k -DPPs on well-conditioned kernels, with experiments supporting our theoretical result. An important direction of future work is improving the sampling time complexity of Partition-DPPs to run in time polynomial in m, k as well as p . Another direction is to explore more constrained-DPPs and to provide efficient sampling algorithms for them.

References

- [1] Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 333–342, New York, NY, USA, 2011. ACM.
- [2] Raja Hafiz Affandi, Emily B. Fox, Ryan P. Adams, and Ben Taskar. Learning the parameters of determinantal point process kernels. *ICML*, 2014.
- [3] M. Artin. *Algebra*. Prentice-Hall, 1991.
- [4] Kobus Barnard, Pinar Duygulu, David A. Forsyth, Nando de Freitas, David M. Blei, and Michael I. Jordan. Matching words and pictures. *Journal of Machine Learning Research*, 3:1107–1135, 2003.
- [5] Jeff Bilmes. Personal communication. 2016.
- [6] Alexei Borodin. Determinantal point processes. *arXiv preprint arXiv:0911.1153*, 2009.
- [7] Alexei Borodin and E.M. Rains. Eynard-Mehta theorem, Schur process, and their pfaffian analogs. *Journal of statistical physics*, 121(3-4):291–317, 2005.
- [8] Richard Brualdi and Hans Schneider. Determinantal identities: Gauss, schur, cauchy, sylvestre, kronecker, jacobi, binet, laplace, muir, and cayley. *Linear Algebra and its Applications*, 1983.
- [9] Niv Buchbinder, Moran Feldman, Joseph Naor, and Roy Schwartz. A tight linear time (1/2)-approximation for unconstrained submodular maximization. *FOCS*, pages 649–658, 2012.
- [10] Ali Çivril and Malik Magdon-Ismail. On selecting a maximum volume sub-matrix of a matrix and related problems. *Theor. Comput. Sci.*, 410(47-49):4801–4811, 2009.
- [11] Ali Çivril and Malik Magdon-Ismail. On selecting a maximum volume sub-matrix of a matrix and related problems. *Theor. Comput. Sci.*, 410(47-49):4801–4811, 2009.
- [12] Ali Çivril and Malik Magdon-Ismail. Exponential inapproximability of selecting a maximum volume sub-matrix. *Algorithmica*, 65(1):159–176, 2013.
- [13] Kenneth L. Clarkson. Coresets, sparse greedy approximation, and the frank-wolfe algorithm. *ACM Trans. Algorithms*, 6(4):63:1–63:30, September 2010.
- [14] Amit Deshpande and Luis Rademacher. Efficient volume sampling for row/column subset selection. *FOCS*, 2010.
- [15] Amit Deshpande, Luis Rademacher, Santosh Vempala, and Grant Wang. Matrix approximation and projective clustering via volume sampling. *Theory of Computing*, 2006.
- [16] J. Edmonds and D. R. Fulkerson. Transversals and matroid partition. *J. Res. Natl. Bur. Stand., Sec. B: Math. & Math. Phys.*, 69B(3):147–153, 1965.
- [17] A.S. Georgiades, P.N. Belhumeur, and D.J. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 23(6):643–660, 2001.
- [18] Jennifer Gillenwater, Alex Kulesza, and Ben Taskar. Near-Optimal MAP Inference for Determinantal Point Processes. *NIPS*, 2012.
- [19] S. A. Goreinov, N. L. Zamarashkin, and E. E. Tyrtyshnikov. Pseudo-skeleton approximations by matrices of maximal volume. *Mathematical Notes*, 62(4):515–519, 1997.
- [20] J. Ben Hough, Manjunath Krishnapur, Yuval Peres, and Bálint Virág. Determinantal processes and independence. *Probab. Surveys*, 3:206–229, 2006.
- [21] Rishabh K. Iyer and Jeff A. Bilmes. Submodular point processes with applications to machine learning. In *AISTATS*, 2015.
- [22] Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies. *J. Mach. Learn. Res.*, 9:235–284, 2008.
- [23] Alex Kulesza and Ben Taskar. k-dpps: Fixed-size determinantal point processes. In *ICML*, 2011.
- [24] Alex Kulesza and Ben Taskar. Determinantal Point Processes for Machine Learning. *Found. Trends® Mach. Learn.*, (2-3):123–286, 2012.
- [25] K.C. Lee, J. Ho, and D. Kriegman. Acquiring linear subspaces for face recognition under variable lighting. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 27(5):684–698, 2005.
- [26] Hui Lin and Jeff Bilmes. Learning mixtures of submodular shells with application to document summarization. *UAI*, 2012.
- [27] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 2004.
- [28] Russell Lyons. Determinantal probability measures. *Publications Mathématiques de l’Institut des Hautes Études Scientifiques*, 98(1):167–212, 2003.
- [29] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical Programming*, 14(1):265–294, 1978.

- [30] Aleksandar Nikolov. Randomized rounding for the largest simplex problem. In *STOC*, pages 861–870, 2015.
- [31] Aleksandar Nikolov and Mohit Singh. Maximizing determinants under partition constraints. In *STOC*, 2016.
- [32] Tomoyuki Shirai and Yoichiro Takahashi. Random point fields associated with certain fredholm determinants i: fermion, poisson and boson point processes. *Journal of Functional Analysis*, 205(2):414 – 463, 2003.
- [33] Marco Di Summa, Friedrich Eisenbrand, Yuri Faenza, and Carsten Moldenhauer. On largest volume simplices and sub-determinants. In *SODA*, 2015.
- [34] L.N. Trefethen and D. Bau. *Numerical Linear Algebra*. 1990.
- [35] L. G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8(2):189–201, 1979.

APPENDIX

Complexity Theoretic Barrier for Exact DPP Sampling under General Matroid Constraints

We show that DPP under transversal matroid constraints can simulate random matchings in bipartite graphs. Transversal matroids are defined using a collection of subsets $S_1, S_2, \dots, S_k \subseteq [m]$ that are not necessarily disjoint. A subset S satisfies transversal matroid constraint if $|S \cap S_i| \leq 1$, for all $1 \leq i \leq k$. See Edmonds and Fulkerson[16] for details.

We can define an underlying bipartite graph (U, V, E) where the left vertices are $U = \{S_1, S_2, \dots, S_k\}$ and the right vertices $V = \{1, 2, \dots, m\}$ and we put an edge between S_i and j , if $j \in S_i$. Suppose we relax the disjointness condition in our partition constraints and use the constraints $|S \cap S_i| = 1$, for $1 \leq i \leq k$, instead. Then even for $K = I \in \mathbb{R}^{m \times m}$, the corresponding constrained DPP gives a uniform distribution over random perfect matchings in the above bipartite graph. It is well-known that the partition function for this is equivalent to permanent which is #P-hard to compute by Valiant’s theorem [35].

Similar to determinantal or Fermion point processes, one can define permanental or Boson point processes. Aaronson-Arkhipov [1] have shown that exact sampling from Boson point processes imply that the polynomial hierarchy collapses to its third level.

Partition-DPPs

We first provide an example demonstrating that the formula for partition function by [21] does not work for partition matroids. Consider a partition $\mathcal{P}_1 = \{1, 2\}$ and $\mathcal{P}_2 = \{3, 4\}$ of $\{1, 2, 3, 4\}$ with $k_1 = k_2 = 1$ and let $f(S) = \det(K_{S,S})$. The formula in Section 3.6 in [21] gives $Z^1(\emptyset, \mathcal{P}_1) + Z^1(\emptyset, \mathcal{P}_2) = [f(1) + f(2)] + [f(3) + f(4)]$, whereas the partition function for bases of this partition matroid will be $f(\{1, 3\}) + f(\{1, 4\}) + f(\{2, 3\}) + f(\{2, 4\})$. This is why sampling from a Partition-DPP is non-trivial and necessitates new algorithms.

We now present proofs our main results. We will first need a useful theorem known as Schur’s identity [8].

Lemma 4.1. (*Schur’s Identity*) Let $A \in \mathbb{R}^{m \times n}$ and $S, T \subseteq [m]$, $S \cap T = \emptyset$ and $B = A - \pi_S(A)$. Then,

$$\det(A_{S \cup T} A_{S \cup T}^T) = \det(A_S A_S^T) \det(B_T B_T^T)$$

We now prove Lemma 3.1 of the main paper.

Proof. (of Lemma 3.1) First, it is clear that $c_{0, \dots, 0} = \det(AA^T)$. Next, let $B = AA^T$ and $B' = B - \sum_{i=1}^p x_i I_i$ and $\text{Perm}(S)$ be the set of permutations of a set S . The sign of a permutation τ , denoted by $\text{sgn}(\tau)$, is equal to 1 if it is a product of an even number of transpositions and -1 otherwise.

$$\det(B - \sum_{i=1}^p x_i I_i) = \det(B') = \sum_{\tau \in \text{Perm}([m])} \text{sgn}(\tau) B'_{1, \tau(1)} B'_{2, \tau(2)} \cdots B'_{m, \tau(m)}$$

The term $c_{m_1-k_1, \dots, m_p-k_p} x_1^{m_1-k_1} \dots x_p^{m_p-k_p}$ is the sum over those permutations τ which fix some set $S \subseteq [m]$ of size $(m-k)$, where $k = \sum_{l=1}^p k_l$ such that for all $1 \leq l \leq p$, we have $|S \cap \mathcal{P}_l| = m_l - k_l$ and the elements $\prod_{i \in S} B'_{i,i}$ contribute $(-1)^{m-k} x_1^{m_1-k_1} \dots x_p^{m_p-k_p}$ and the co-efficient comes from the constant term in

$$\sum_{\tau \in \text{Perm}([m]-S)} \text{sgn}(\tau) \prod_{i \notin S} B'_{i, \tau(i)}.$$

The above follows just from the Leibniz formula for the determinant [34] and then collecting the co-efficients of each monomial.

Each term in this sum is the c_0 term of a principal minor of B satisfying the partition constraints. Thus, the sum is equal to

$$\begin{aligned} & \sum_{\substack{S \subseteq [m] \\ \forall i \in [p] : |S \cap \mathcal{P}_i| = m_i - k_i}} \det(B_{[m]-S, [m]-S}) \\ &= \sum_{\substack{S \subseteq [m] \\ \forall i \in [p] : |S \cap \mathcal{P}_i| = k_i}} \det(B_{S,S}) \\ &= \sum_{\substack{S \subseteq [m] \\ \forall i \in [p] : |S \cap \mathcal{P}_i| = k_i}} \det(A_S A_S^T) \end{aligned}$$

Hence,

$$c_{m_1-k_1, \dots, m_p-k_p} = (-1)^{m-k} \sum_{\substack{S \subseteq [m] \\ \forall i \in [p] : |S \cap \mathcal{P}_i| = k_i}} \det(A_S A_S^T)$$

□

Finally, we prove our main theorem which is Theorem 3.2 in the main paper.

Proof. (of Theorem 3.2) Define \mathcal{C}_1 be a set of constraints as follows

$$\mathcal{C}_1 = \left\{ T \subseteq [m] \mid |T| = k - t - 1, \forall l \in [p] : |(S \cup \{i\} \cup T) \cap \mathcal{P}_l| = k_l \right\}$$

In other words,

$$\mathcal{C}_1 = \left\{ T \subseteq [m] \mid |T| = k - t - 1, |T \cap \mathcal{P}_j| = k_j - t_j - 1, \forall l \in [p] \setminus \{j\} : |T \cap \mathcal{P}_l| = k_l - t_l \right\}$$

We, once again, stress that the above follows since the partitions are considered to be disjoint. Similarly define \mathcal{C}_2 as

$$\mathcal{C}_2 = \left\{ T \subseteq [m] \mid |T| = k - t, \forall l \in [p] : |T \cap \mathcal{P}_l| = k_l - t_l \right\}$$

Then,

$$\begin{aligned}
& \Pr(X_{t+1} = i | X_1 = i_1, \dots, X_t = i_t) \\
&= \frac{\sum_{\substack{(i_{t+2}, \dots, i_k) \in [m]^{k-t-1} \\ \forall l \in [p] : |(S \cup \{i\} \cup \{i_{t+2}, \dots, i_k\}) \cap \mathcal{P}_l| = k_l}} \Pr(i_1, \dots, i_t, i, i_{t+2}, \dots, i_k)}{\sum_{\substack{(i_{t+1}, \dots, i_k) \in [m]^{k-t} \\ \forall l \in [p] : |(S \cup \{i_{t+1}, \dots, i_k\}) \cap \mathcal{P}_l| = k_l}} \Pr(i_1, \dots, i_t, i_{t+1}, i_{t+2}, \dots, i_k)} \\
&= \frac{(k-t-1)! \sum_{\mathcal{C}_1} \det(A_{S \cup \{i\} \cup T} A_{S \cup \{i\} \cup T}^T)}{(k-t)! \sum_{\mathcal{C}_2} \det(A_{S \cup T} A_{S \cup T}^T)} \\
&= \frac{\sum_{\mathcal{C}_1} \det(A_{S \cup \{i\} \cup T} A_{S \cup \{i\} \cup T}^T)}{(k-t) \sum_{\mathcal{C}_2} \det(A_{S \cup T} A_{S \cup T}^T)} \\
&= \frac{\sum_{\mathcal{C}_1} \det(A_S A_S^T) \det(B_{\{i\} \cup T} B_{\{i\} \cup T}^T)}{(k-t) \sum_{\mathcal{C}_2} \det(A_S A_S^T) \det(B_T B_T^T)} \quad \text{by Lemma 4.1} \\
&= \frac{\sum_{\mathcal{C}_1} \|b_i\|^2 \det((C_i)_T (C_i)_T^T)}{(k-t) \sum_{\mathcal{C}_2} \det(B_T B_T^T)} \quad \text{by Lemma 4.1 applied on B} \\
&= \frac{\|b_i\|^2 \sum_{\mathcal{C}_1} \det((C_i)_T (C_i)_T^T)}{(k-t) \sum_{\mathcal{C}_2} \det(B_T B_T^T)} \\
&= \frac{\|b_i\|^2 |c'(C_i C_i^T)|}{(k-t) |c''(B B^T)|} \quad \text{by Lemma 3.1}
\end{aligned}$$

□

MAP Inference for k -DPPs

Firstly, denote $d(x, E)$ to be the distance of a vector x to the subspace E , i.e., it is the l_2 -norm of the projection of the x onto the subspace orthogonal to E .

Proof. (of Theorem 4.1) A simple greedy algorithm that starts with $S = \emptyset$ and, in k steps, adds one element at a time that maximizes $\det(A_{S \cup \{i\}} A_{S \cup \{i\}}^T)$ and updates S to $S \cup \{i\}$, gives a $2^{\mathcal{O}(k \log k)}$ factor approximation to the MAP inference problem for k -DPPs [11]. We can now run a local search where we pick an element $i \in S$ and replace it with some $j \notin S$ that gives maximum improvement as long as

$$\det(A_{S \setminus \{i\} \cup \{j\}} A_{S \setminus \{i\} \cup \{j\}}^T) > \left(1 + \frac{\epsilon}{k}\right) \det(A_S A_S^T),$$

and update S to $S \setminus \{i\} \cup \{j\}$. Since we start with a greedy initialization that gives $2^{\mathcal{O}(k \log k)}$ approximation, the local improvement must halt in $\mathcal{O}(k^2 \log k / \epsilon)$ steps. When the local improvement stops, for all $i \in S$ and $j \notin S$, we have

$$\det(A_{S \setminus \{i\} \cup \{j\}} A_{S \setminus \{i\} \cup \{j\}}^T) < \left(1 + \frac{\epsilon}{k}\right) \det(A_S A_S^T),$$

This implies that, for all $i \in S$ and $j \notin S$,

$$d(a_j, \text{span}(S \setminus \{i\}))^2 < \left(1 + \frac{\epsilon}{k}\right) d(a_i, \text{span}(S \setminus \{i\}))^2.$$

Therefore, summing the LHS over all $j \notin S$ and i we get

$$\begin{aligned} & \|A - \pi_{S \setminus \{i\}}(A)\|_F^2 \\ & < \left(1 + \frac{\epsilon}{k}\right) (n - k) d(a_i, \text{span}(S \setminus \{i\}))^2, \end{aligned}$$

for all $i \in S$. For any $i \in S$, the LHS is at least the best rank- $(k - 1)$ approximation error for A under the Frobenius norm, which is $\sum_{t \geq k} \lambda_t(AA^T)$. Taking the product of the RHS over $i \in S$ and upper bounding $\prod_{i \in S} d(a_i, \text{span}(S \setminus \{i\}))^2$ by $\det(A_S A_S^T)$, we obtain

$$\frac{1}{\left(1 + \frac{\epsilon}{k}\right)^k} \left(\frac{1}{n - k} \sum_{t \geq k} \lambda_t(AA^T) \right)^k < \det(A_S A_S^T),$$

Suppose we do not have a $2^{k \log \kappa(A) + 1}$ or $2\kappa(A)^k$ -approximation when the local search stops, then $\det(A_S A_S^T) < \frac{1}{2} \kappa(A)^{-k} \max_{|S|=k} \det(A_S A_S^T) \leq \frac{1}{2} \kappa(A)^{-k} (\lambda_1(AA^T))^k$, which is at most

$$\frac{1}{2} \left(\frac{1}{n - k} \sum_{t \geq k} \lambda_t(AA^T) \right)^k, \quad \text{giving a contradiction.}$$

□

Experiments

Corel16k Task for “Limitations of k -DPPs”

In this section, we describe the data and experimental results of the multi-label image search task which demonstrates another limitation of k -DPPs.

Data and DPP Kernel

For the second task, we use the corel16k dataset [4]. The images were cropped to consider them to be of the same size and centered. We fix our label space to be a subset of size 3 in two experiments. In the first experiment, the label space is {sky, sun, tree}. We took 9 images (numbered 1 to 9) where they have one or more labels from this label set and Table 1 shows the images in each label/partition. Each image had at most 4 labels and those images were chosen which did not have more than 2 labels not from this label set. In the second experiment, the label space considered was {prop, plane, formation} and again 9 images were considered. Since the corel16k dataset does not provide the actual images, we only check for duplicates in the set of images sampled. We just used the 46 features provided in the dataset itself. The kernel is just the dot-product of the features of both images.

| Label | Corresponding Images |
|-------|----------------------|
| Sky | 1,4,6,7,8 |
| Sun | 1,2,3,4,5,6,7 |
| Tree | 5,6,7,8,9 |

Table 4: Labels and Corresponding Images for Experiment 1 of the Corel16k task

Results

This task was performed to illustrate the limitation of independent k_i -DPPs which is the possible duplication in the selected images. Clearly, k -DPPs and Partition-DPPs will not choose the same image multiple times, however, independent k_i -DPPs may suffer from this. This is especially likely

in cases where some elements, which occur in multiple partitions, are more likely to be sampled due to the kernel promoting them. The two tables in Table 3 provide the list of sampled sets (each column denotes a partition) in three different runs of the experiment for a given configuration of k_i s. The first three rows in both correspond to the k_i configuration of (3, 3, 3) and the bottom three correspond to (2, 3, 2). The images were labelled from 1 to 9. As can be seen from the table, there is considerable duplication of elements in both the cases.

| Sky | Sun | Tree | Prop | Plane | Formation |
|-------|-------|-------|-------|-------|-----------|
| 1,4,7 | 1,3,6 | 5,6,9 | 1,3,4 | 4,6,9 | 5,6,9 |
| 4,7,8 | 1,2,6 | 5,6,7 | 2,3,6 | 2,4,7 | 4,5,8 |
| 1,6,8 | 2,3,7 | 6,8,9 | 4,6,9 | 6,8,9 | 4,5,8 |
| 1,4 | 1,3,6 | 6,7 | 1,2 | 2,7,9 | 4,8 |
| 4,6 | 5,6,7 | 8,9 | 4,9 | 3,4,8 | 5,6 |
| 7,8 | 2,3,7 | 5,7 | 1,3 | 2,3,9 | 6,9 |

Table 5: Elements sampled from the two different sets of the independent k_i -DPPs to illustrate duplication in chosen elements

MAP Inference for k -DPPs

We first present the results of the first synthetic task which consists of 50 points in 40 dimensions and each $A_{i,j} \sim 5 * \text{Uniform}(0,1)$. The results are provided in Table 3.

| k | Method | Value |
|----------|------------|---------|
| $k = 4$ | Greedy | 1.67e9 |
| | Our method | 1.67e9 |
| | Optimal | 2.05e9 |
| $k = 10$ | Greedy | 4.27e16 |
| | Our method | 4.27e16 |
| $k = 20$ | Greedy | 3.09e31 |
| | Our method | 3.09e31 |
| $k = 25$ | Greedy | 9.93e37 |
| | Our method | 1.82e38 |

| k | Method | Value |
|------------------------|------------------|---------|
| $k = 35$ | Greedy | 2.87e48 |
| | Our method | 3.01e48 |
| $k = 39$ | Greedy | 9.69e50 |
| | Our method | 1.06e51 |
| | Symmetric Greedy | 1.59e50 |
| $k = 40 = \text{rank}$ | Greedy | 2.07e51 |
| | Our method | 2.91e51 |
| | Softmax | 1.24e51 |

Table 6: Results of the first synthetic dataset for MAP Inference

Our second synthetic dataset consists of 200 points in 200 dimensions and each $A_{i,j} \sim \text{Uniform}(0,1)$. In this case, Greedy, our algorithm and Symmetric Greedy all picked all 200 of the elements. The results are shown in Table 4.

| k | Method | Value |
|----------|------------|---------|
| $k = 5$ | Greedy | 3.62e7 |
| | Our method | 3.89e7 |
| $k = 10$ | Greedy | 1.30e14 |
| | Our method | 1.48e14 |
| $k = 20$ | Greedy | 5.34e26 |
| | Our method | 5.99e26 |

| k | Method | Value |
|-----------|------------|----------|
| $k = 50$ | Greedy | 2.96e62 |
| | Our method | 3.18e62 |
| $k = 100$ | Greedy | 5.14e98 |
| | Our method | 5.61e98 |
| $k = 150$ | Greedy | 1.03e155 |
| | Our method | 1.08e155 |

Table 7: Results of the second synthetic dataset for MAP Inference

For this dataset, the kernel is very well-conditioned and our algorithm beats greedy very early on.

We finally have the results for the matched summarization real-world task of [18] in Table 5.

| k | Method | Value | k | Method | Value |
|-----------|------------|---------|-----------|------------------|---------|
| $k = 5$ | Greedy | 2.95e3 | $k = 250$ | Greedy | 1.03e78 |
| | Our method | 2.95e3 | | Our method | 1.11e78 |
| $k = 10$ | Greedy | 9.55e5 | $k = 300$ | Greedy | 2.79e81 |
| | Our method | 1.01e6 | | Our method | 3.15e81 |
| $k = 25$ | Greedy | 8.07e12 | $k = 312$ | Greedy | 2.46e81 |
| | Our method | 8.54e12 | | Symmetric Greedy | 6.63e80 |
| $k = 50$ | Greedy | 2.58e23 | $k = 327$ | Our method | 2.82e81 |
| | Our method | 2.83e23 | | Greedy | 5.78e80 |
| $k = 100$ | Greedy | 1.25e42 | $k = 350$ | Softmax | 6.05e80 |
| | Our method | 1.32e42 | | Our method | 4.63e80 |
| $k = 150$ | Greedy | 4.77e57 | $k = 380$ | Greedy | 2.78e78 |
| | Our method | 4.89e57 | | Our method | 2.94e78 |
| $k = 200$ | Greedy | 6.26e69 | $k = 380$ | Greedy | 3.03e72 |
| | Our method | 7.09e69 | | Our method | 3.22e72 |

Table 8: Results of the real-world dataset for MAP Inference